The Smart Floor: A Mechanism for Natural User Identification and Tracking

Robert J. Orr and Gregory D. Abowd
Graphics, Visualization, and Usability (GVU) Center
Georgia Institute of Technology
Atlanta, GA 30332-0280 USA
+1 404 894 5103
{rjo, abowd}@cc.gatech.edu

ABSTRACT
We have created a system for identifying people based on their footstep force profiles and have tested its accuracy against a large pool of footstep data. This floor system may be used to transparently identify users in their everyday living and working environments. We have created user footstep models based on footstep profile features and have been able to achieve a recognition rate of 93% using this feature-based approach. We have also shown that the effect of footwear is negligible on recognition accuracy.

Keywords
Interaction technology, ubiquitous computing, tactile I/O, user identification, biometrics, indoor positioning, intelligent home environments, intelligent systems.

INTRODUCTION

Transparent user identification has in recent years become a goal of computer science researchers. With the advent of ubiquitous computing [17] and smart environments, transparent user identification has become an even more pressing goal than before the rise of these paradigms. If a computer or environment could transparently identify the user, it could customize its interface and behavior to match the preferences, history, and context of that particular user. Computers could become easier to use, and could themselves become more transparent overall. But the system must provide a service or services of enough value that the user will tolerate the additional technology, no matter how transparent it may be.

For example, many users may appreciate a system that tracks certain objects around their living spaces for them. Frequently lost objects, such as keys, wallets, and glasses, could be tagged with small radio frequency ID tags, and their location could be correlated with the location of people in the house [8]. As an example, if Mary were to walk out of the house with a set of keys and Joe needed to locate the keys some time later, the system could inform Joe that “The keys were last seen 30 minutes ago at the front door with Mary.” Joe could then deduce that the keys were with Mary and coordinate with her. In this case, the identity and location of Mary is an important piece of information. Furthermore, if the system cannot transparently track and identify users, they will be much less likely to use the system, and the services the system offers will be much less likely to be successful. On the other hand, if the system is transparent enough, the ability to track frequently lost objects may be a compelling enough service that users are willing to be tracked and identified, and willing to have that information made available to a small group of people within the house.

However, most identification methods to date have not been especially transparent. The Active Badge system [16] is one of the most widely used systems and illustrates some of the problems with many identification schemes. (Radio frequency identification systems, or RFID, have many of the same features and problems.) First and foremost, the user must carry a badge or tag in order to be identified. While this can be a feature at times--if the user desires privacy, all she does is remove the badge--it can be an impediment to use and also narrowly defines the environments in which the system can be used. For example, the Active Badge system is not particularly amenable to use in a smart house: users will not wear the badge while sleeping (in order, for example, that the house can identify them when they arise to use the facilities in the middle of the night), or while doing work in the yard. They must remember to put it on when they arise or come back into the house. In addition, adding new users, such as frequent visitors, requires another physical badge or tag. Finally, badge systems only provide gross positioning. The best badge-based indoor positioning system to date [18] only has a resolution of 6 feet. In many cases, we will want to know the position of a user to a finer resolution.

There has also been much work recently that has focused on more passive forms of user recognition, such as face recognition using video and voice recognition using audio. These types of recognition do not require that the user carry any tag or badge; they utilize only biometric data from the user. Video and audio can also both be used to track the location of users in a space, and to a much finer resolution than badge and tag systems. However, these technologies have problems, too. Video recognition is stymied by occlusions, shadows, and lighting inconsistencies, and won't work at all in the dark. Audio recognition suffers from problems of background noise and requires the user to speak in order to function (not very useful in the middle of the night when one is trying to tiptoe quietly without waking the other occupants of the house).

An Alternative Biometric Approach

Passive biometric approaches have the advantage that the user does not need to carry anything or remember anything. The badge and tag systems require the user to carry a badge or tag, but they work just fine in noisy environments and occluded or dark rooms. We have designed a system that has many of the advantages and few of the drawbacks of both classes of systems. Our Smart Floor system measures the force exerted on a floor tile and is able to recognize the user based on their footstep force profile as they walk over the tile.

The floor has a number of characteristics that make it an obvious choice for instrumentation: users always walk over it; it is always there (even in the dark); and it can sense information not only about users but also about objects. With the Smart Floor, the user does not need to carry anything (like a badge) or remember anything (like a password); she simply walks over the floor tile and the system utilizes the user's biometric data for recognition. The Smart Floor also works fine when the room is dark or noisy, and it does not care if a view of the user is occluded. In addition, by its very nature the floor gives accurate position information. Lastly, the algorithms for identification and tracking are fairly simple and not computationally intensive.

Purpose of the Project

The purpose of the Smart Floor project [7] has been to create and validate biometric user identification based on users' footsteps. As mentioned above, we have outfitted a floor tile with force measuring instruments and are using the data gathered as users walk over the tile to identify them. We rely on the fact that footstep profiles are unique enough within a small enough group of people to provide recognition capabilities matching or exceeding the capabilities of other biometric technologies. (We will address this claim further below.) Specifically, we have been able to achieve a 93% overall user recognition rate with our system, and have been able to show that footwear is not a significant factor in identifying users. Furthermore, we have created a system that can transparently identify users and now allows us to prototype useful services for users.

We had a number of research goals for the Smart Floor system at the start of this project:

- Create an accurate system for recognizing a user's identity from their footsteps;
- Investigate the similarity of users' footsteps and show that for a small group of users (up to about 15), different users' footstep profiles are dissimilar enough for our system to work satisfactorily;
- Create a system that can track a user over an area larger than just a single floor tile;
- Use the system in a real environment with real users and real applications.

In this paper, we will describe the progress we have made towards the first three of these four goals, and our plans for the fourth goal.

Technology Tradeoffs

It is important to note at the outset that we do not intend the Smart Floor to be a single technology replacement for the other types of identification technologies. If that were our goal, we would have aimed to design a system that gives perfect recognition and is transparent to use (i.e., the user need not expend any additional effort for the technology to do its job). Rather, we intend that the Smart Floor will work in conjunction with other technologies. For example, the floor system may provide a set of weighted identities to a voting system that has inputs from other systems such as voice and face recognition systems. Further, a video recognition system may overcome some of the shortcomings of the floor system (at least in its current form), such as the inability to distinguish between people who have very similar walking profiles but who may be of widely different heights.

No identification system is perfect. A technology may give very accurate (or even perfect) results, but these technologies are not usually transparent. Transparent systems usually do not give 100% accurate results. Our floor system falls into this later category and this is why we intend that it be used with other recognition systems.

EXPERIMENTAL SETUP

In this section, we will start with an short explanation of ground reaction force, and then detail the equipment and methods we have used to gather the footstep profiles, model each user, and compare unidentified footsteps to the known user models.

Ground Reaction Force

In the biomechanics literature, the reaction that a measuring device produces in response to the weight and inertia of a body in contact with that device is called *ground reaction force* (GRF) [2,11,15]. In our case, we are measuring the GRF of the walker's foot as he or she walks over out measuring tile. A sample GRF profile is shown below in Figure 1; GRF is represented on the vertical axis, time is represented on the horizontal axis.

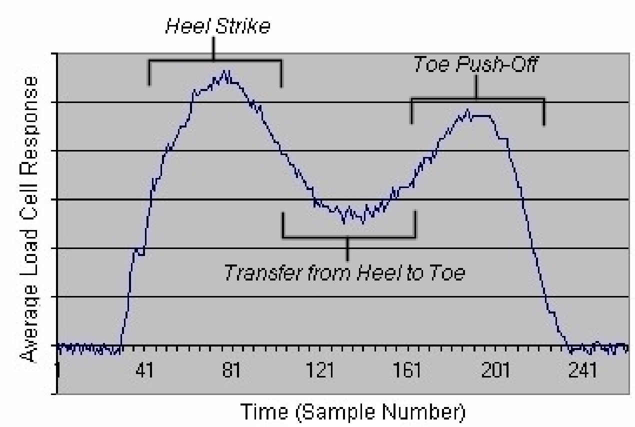

Figure 1. *Sample ground reaction force (GRF) profile.*

The heel strike is represented by the left hump in the figure, while the right hump represents the force exerted by the toe push-off as the foot leaves the tile. The middle section of the curve shows the transfer of weight from the heel to the toe.

In our system, we are measuring only the vertical component of ground reaction force. More sophisticated systems are additionally able to measure both horizontal components, as well as torsional components. These GRF components may be useful in identifying a user (they may indicate pronation or supination, for example). However, the additional cost and complexity required to capture these aspects of the GRF make doing so prohibitive.

Hardware

The hardware we have used to gather GRF profiles consists of three components: load cells, a steel plate, and data acquisition hardware.

Load Cells and Load Plate

To reliably gather GRF profiles, we have created a floor tile consisting of a 3/8” thick steel plate mounted on four industrial load cells; the load cells are laid on the floor. This load plate measures 50cm by 50cm. The load cells are 25mm diameter cylinders with 3mm high, 9.7mm wide load buttons on the top surface (see Figure 2). Each load cell has a capacity of 500 pounds, giving the plate a total capacity of 2000 pounds. The steel plate was machined with a 10mm wide, 1.5mm deep hole in each corner so that the load buttons fit snugly into the holes. This arrangement keeps the load plate from moving horizontally when horizontal force is applied to it.

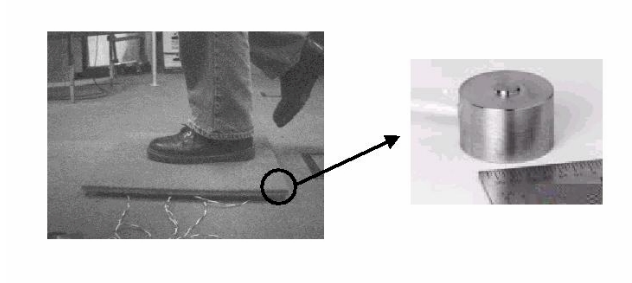

Figure 2. *Smart Floor plate (left) and load cell (right).*

Floor Setup

We have placed a 50cm square tile of short-pile industrial carpeting on top of the load plate in order to keep users from slipping as they walk across it. This assembly rises about 1” above the floor. We currently have plans to install this load plate into a smart home and the plate will be mounted so it is flush with the floor surface around it.

Data Acquisition

The load cells require a 10-15 Vdc excitation voltage, and return 2.2 mV of output per Volt of excitation. For example, if a 200 pound load is placed on a 500 pound capacity load cell with a 10 V excitation voltage, it will have a 8.8mV output. The excitation voltage is provided by an external 12V power supply, and the resulting signal is conditioned by a National Instruments SC-2043-SG board. The conditioning board converts its input signal, which is a floating, non-referenced signal (each load cell is composed of four strain gauges arranged in a Wheatstone bridge), to a signal that is referenced to ground. The signals from all four load cells are then converted from analog to digital by a National Instruments PCI-1200 data acquisitions board in a standard Pentium PC. The system runs under the Windows NT operating system.

Software

The software we have created for the smart floor falls into two categories: data acquisition and user modeling. The data acquisition software is used to acquire individual footsteps from the load plate, while the user modeling software is used to create models of each individual user’s GRF profiles and then compare unknown identity footsteps with previously created models.

Data Acquisition

We have written a utility to gather the footstep signals from the data acquisition hardware, average and calibrate the signals from the four load cells, automatically segment out each individual footstep profile, and store the results to disk. The software uses a set of software drivers that National Instruments provides to communicate with the data acquisition hardware. It has been shown that most of the energy in footsteps lies below 250 Hz [1], and we therefore decided to sample each load cell 300 times per second for a resolution limit of 150 Hz. We have found this sampling rate to be more than adequate for the modeling methods we have used; the features we use do not depend on the higher frequency components of the signal. A sample GRF profile was shown above in Figure 1; it shows the average of all four load cells.

Our software adjusts the signal from each load cell using additive and multiplicative calibration constants provided by the manufacturer. We also subtract the mean force of the load plate (the first 500 samples are gathered for this in order to calculate unloaded mean and variance) and account for variations in the excitation voltage (we acquire the excitation voltage for this purpose). The software then averages the signals together and segments each footstep profile from the data stream; each footstep profile is written to disk. However, as the interface in Figure 3 shows, we also have the ability to save all five channels of raw data (four load cells and the value of the excitation voltage). The software can be easily set to gather a specific number of footsteps, and will provide feedback when it reaches that number.

GRF Profile Features

In modeling each user's footsteps, we have chosen ten footstep profile features to use as markers for each GRF profile; they are show below in Figure 3. For our features we are using:

- the mean value of the profile,
- the standard deviation of the profile,
- the length of the profile (the number of samples),
- the total area under the profile curve,
- the coordinates of the maximum point in the first half of the profile (x_{max1} and y_{max1}),
- the coordinates of the maximum point in the last half of the profile (x_{min} and y_{min}),
- and the coordinates of the minimum point between the two maximum points (x_{max2} and y_{max2}).

There are other features that can be used in this approach, such as average curve slope between onset and first maximum, but we have chosen to limit our feature set to those mentioned above simply because we felt that these ten features were already probably more than we needed to sufficiently characterize the footstep profiles. Note also that not all of the features we have used are orthogonal to the others. For example, the mean is trivially calculated from the area and the length.

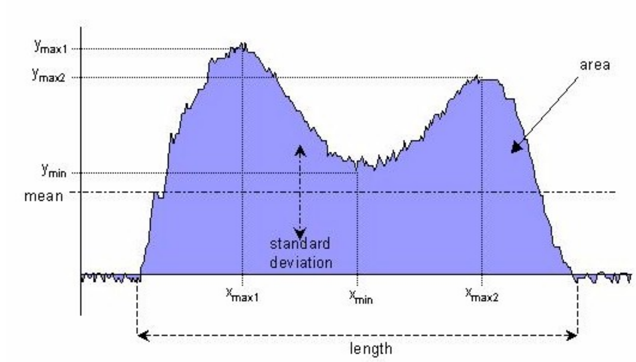

Figure 3. *Footstep profile features.*

Nearest-Neighbor

In the feature classification approach, we have used the simple approach of placing each footstep into a ten-dimensional feature space, and using a nearest neighbor search to identify unknown footsteps. Each feature listed above is one dimension in this ten dimensional space. We calculate the Euclidean distance in feature space from the unknown footstep to each known footstep and the identity of the closest match can be chosen as the identity of the unknown footstep. Since the scales and units of the features is different, each dimension of the feature vector must be scaled by a constant so that each dimension can be compared with the others. This way the largest valued feature will not dominate the other features in the distance calculation. More precisely, let an arbitrary footstep instance x be described by the feature vector $\langle a_1(x), a_2(x), \ldots a_n(x) \rangle$, where $a_r(x)$ denotes the value of the r^{th} attribute of instance x (e.g., mean, length, etc.). Then the distance between two footstep instances x_i and x_j is defined to be $d(x_i,x_j)$, where

$$d(x_i, x_j) \equiv \sqrt{\sum_{r=1}^{n} \left(\frac{a_r(x_i) - a_r(x_j)}{\max(a_r)} \right)^2} \qquad \text{(Eq. 1)}$$

In Equation 1, r ranges over all the feature types in the feature set (n in our case is 10), and we take x_i to be the unknown footstep instance and x_j to be the known footstep instance. Max(a_r) gives us the maximum value that that particular feature takes over all the footsteps in our stored training data set; it serves as the calibration constant.

We calculate this distance between the unknown identity footstep and every known identity footstep in our data set. The simplest way to choose the identity of the unknown footstep is to use the identity of the footstep that is closest in feature space. A more sophisticated method is based on footstep clustering. Each user has given us a number of known training footstep samples; these samples form clusters of points in feature space (we will show this below). The distance is computed from the unknown identity footstep to every footstep in every cluster. The identity of the cluster with the lowest average distance is chosen as the identity of the unknown footstep.

This feature classification approach contains a few assumptions. First, we assume that the GRF profiles can be accurately and sufficiently modeled by a small number of well-chosen features. If the features we have chosen are insufficiently descriptive or too many of the features are non-orthogonal, this method will be unable to separate the clusters of footsteps in feature space and the recognition accuracy will suffer. Another assumption is that the clusters of footsteps formed by a single user are inherently separable from the other clusters. If this were not so, no choice of features, no matter how optimal, could adequately distinguish between different users. We will show below that the cluster are sufficiently separable given the features we have chosen.

RESULTS

We will now examine our experimental results.

Experimental Population and Conditions

We gathered GRF profiles from 15 subjects, 12 male and 3 female. For each subject, we gathered separate data for left and right feet. We attempted to gather data for as many shoe types as possible. For two subjects, we had data for four different pairs of shoes; for five subjects, we had data for three pairs of shoes; three subjects gave data for two pairs of shoes; and the remaining five subjects only gave data for one pair of shoes. For subjects who gave data for more than one shoe pair, one of the conditions for which we gathered data was barefoot. Subjects were instructed to walk over the Smart Floor tile in a normal gait, and to place their foot in roughly the center of the tile. Each footstep was taken with the subject walking in the same direction over the tile.

We refer to one foot and one shoe type as a single condition. E.g., "Joe's left foot while he was wearing tennis shoes" was a single condition. We gathered 20 footsteps per condition, half of which we would use for training and half for testing. In total, we gathered 1680 footstep profiles.

Clustering of Footfall Profiles in State Space

We had two related hypotheses for the GRF profiles we gathered. The first was that the footsteps for a single condition would be clustered closer together in feature space than they would be to the points in any other cluster. In other words, a person's right footstep (for example) while wearing a particular kind of shoe is more similar to his other right footsteps wearing that same shoe than to any other footstep type. Our second hypothesis was that one cluster for a particular person would be closer to the other clusters for that person than to any cluster from any other person. In order to test these two hypotheses, we measured the distance in feature space from every point to every other point and then computed average cluster distances. To calculate distances, we used the approach outlined above.

The results of our work showed that for 48 of the 74 clusters (65%), the smallest distance was to itself (i.e., its average internal distance, or the average distance between the points in the cluster, was smaller than the average distance to other clusters). However, in 17 of the remaining 26 cases, the smallest distance was to another cluster for that same user. For the remaining 9 clusters the smallest distance was to another user's cluster. Our first hypothesis does not stand up strongly, but if we modify it to claim that a cluster will be closer to other clusters for that same user than clusters for other users, it is correct for 65 of the 74 clusters (88%). Ultimately, what we can conclude from these results is that *footwear does not greatly affect the ability of our approach to identify the user by his footsteps*.

In order to fully test this conclusion, we need to examine our second hypothesis, whether clusters for one user are closer to each other than to clusters for another user. We have not completed these analyses as of this writing, but from the first analysis and from the results of our accuracy testing so far, we expect to confirm our hypothesis.

Recognition Accuracy

The best test of our system is to examine the recognition accuracy. As mentioned earlier, we used feature space modeling (with nearest neighbor recognition) to create user footstep models. We used half the gathered data for training the models and the other half for testing them.

In the feature space approach, there are a number of ways to classify an unknown footstep. In first-nearest neighbor, we take the identity of the single closest point in feature space as the identity of the unknown footstep. In this case, we achieved only 16% accuracy; correct identity in this case was correct user, correct foot, and correct shoe. This is due to the fact that the points of the clusters in feature space are interspersed with the points of other clusters, even though the distances between the points in the each cluster are usually less than the distances to the points in another cluster. A point from another may just be closer than the point from the correct cluster.

A better approach is to calculate the distance to all the points in a particular cluster and take the average. When we tested this approach, the recognition rate rose to 36%. Again, a correct match was counted when the system returned the correct user, correct foot, and correct shoe. The overall goal, however, is only to identify someone, to return the correct user only. When we dropped the criterion of correct shoe, the accuracy rose to 75%. And when we only looked at identity, we were able to achieve 93% accuracy. Table 1 below summarizes our results.

Condition	Recognition Accuracy
Correct user, foot, and shoe (first-nearest neighbor)	16%
Correct user, foot, and shoe (cluster averaged)	36%
Correct user and foot	75%
Correct user	93%

Table 1. *Recognition accuracy vs. condition.*

RELATED WORK

Addlesee, et al. [1] designed a system similar to our Smart Floor system. They used load cells to measure force on a floor tile and used footstep data to perform user identification. To do this, they created user footstep models with hidden Markov models (HMMs) and compared unknown identity footsteps against the stored HMMs. To improve the recognition accuracy, they experimented extensively with the HMM parameterization and were able to achieve 91% correct footstep identification in the best case. While this project is similar to our own, it did not investigate the breadth of walking conditions that we have, and it did not investigate methods other than HMMs for modeling user footsteps. We have used some of the results of this work, but we have gone farther in terms of characterizing the similarity of user footsteps and have tested users under far broader conditions.

Pinkston et al. [12] created a dance floor system embedded with strips of force-sensitive resistors (FSRs). The FSRs were used to trigger MIDI sounds and sequences. Paradiso et al. [9] have built a floor system that uses small Doppler radar units and a grid of piezoelectric wires to track a user's position and orientation. The system was used in an audio installation to launch and modify complex sounds and sequences. In another system, Paradiso and Hu [10] instrumented a pair of dance sneakers with piezoelectric pads, force sensitive resistor pads, an accelerometer, and an electronic compass to sense foot impact, pressure, flex, and orientation. They also used sonar mounted around the edge of the performance area to detect the position of the dancer. Each of these systems were designed for use in an artistic setting and concentrated mostly on user position and orientation within a small space. They do not address user identity at all.

FUTURE WORK

There are many issues that we would like to explore with the Smart Floor. Many of these issues relate to using the Smart Floor in a "live use" situation, with real users and real applications, not just in the laboratory. We also have a number of more technical issues regarding the techniques used to model users' GRF profiles.

Live Use in the Aware Home: Applications

The location we are currently planning to deploy the Smart Floor is in the Aware Home, a technologically advanced three-story house we are building at Georgia Tech [4]. The house will contain a number of different sensing technologies, including video- and audio-based technologies, standard environmental sensing (light, temperature, etc.), appliance control and sensing, and location detection. The larger technological goals of the Aware Home are 1) to be able to sense what is happening in the house at any time, including the quantities of who, what, where, and when; and 2) to be able to provide any kind of information to the occupants at any time. The Smart Floor system obviously falls under the first category.

In conjunction with the project architect, we have integrated ten Smart Floor tiles at strategic locations in the Aware Home, including house entrances, kitchens, hallway entrances, and bedroom entrances. These tiles will be suspended in the floor structure and will be flush with the rest of the floor. Each tile location will be able to identify the users as well as provide implicit location information.

One larger goal of the Aware Home project is to support normal users in daily activities. In order to accomplish this, we may target families with children or the elderly, and recruit a family or individuals to live in the Aware Home as their primary residence. This would provide invaluable user experience with which we can evaluate the Smart Floor.

Location Tracking Network

Another component to the Smart Floor system is a location tracking grid. Our goal is to create a location tracking grid that can track a user anywhere in an environment such as the Aware Home. We are currently implementing a grid of piezoelectric wires that will give finer-grained location information than can be derived from the implicit location information from the tiles. Piezoelectric wires generate a voltage potential when compressed or stretched [6], and when arranged in a grid give the location of the compression. As wires have some thickness and may be difficult to attach to the floor surface, we are investigating other technologies such as deformation sensitive fiber optic threads. Another approach to location tracking is to use a network of vibration sensors or audio microphones attached to the underside of the floor surface, and measure signal time-of-flight to triangulate location [19].

In combining the Smart Floor tiles and the location tracking grid, there are a number of issue we will investigate. Tracking multiple users is foremost among these. If two users trigger the same grid lines, for example, it is impossible to detect each user's location.

Also, when two users cross paths, we will need to disambiguate the resulting locations of the users. A technique such as Kalman filtering may help us solve both of these problems: a Kalman filter uses recent movement history to make a prediction about current and future movement. This may an area in which it would be useful to combine the output of our system with that of a video-based tracking system.

Another issue we would like to address is the use of location history as location history and movement patterns may reveal much about behavior and intention. We will address how best to store location history information and integrate this information with the other sensor and context information available from the Aware Home.

Privacy

Privacy will obviously be of major concern with our system (and in the Aware Home in general) and we have devoted, and will continue to devote, attention to this difficult issue. Foremost among our goals here is that the user be informed about what information is being gathered about him and be able to control the gathering of that information and the flow of that information after it has been gathered. In the Aware Home, we will design an "information firewall" into the system in order to ensure that sensitive information is not accessible to the outside world. Beyond that, we need to give occupants control over whether the system will gather footstep and location data. To accomplish this, it would not be difficult to place a standard wall switch near the Smart Floor tile to control power to it and its surrounding location tracking grid (e.g., the switch would turn off all floor sensing in that particular room). We will also need to address how to provide feedback to the user that the system is on or off.

We may also be able to use the data gathered by the system itself to control the system. For example, very high amplitude input, such as a "stomp" footstep, could instruct the system to turn off user recognition until it the next time it sees very high amplitude input. Other types of floor "gestures" may be associated with other system actions.

It is important to note that the Smart Floor system has been designed to be used by a cooperative population of users: the system can be easily fooled by walking over it with a strange gait or by walking around the floor tile altogether. The system is not meant to clandestinely monitor users; system training must happen explicitly. Rather, our goal is to provide useful information to a larger system that can provide users with useful services, motivating them to use the system. However, we still need to provide users the opportunity to turn the system off if they so desire.

Integration with Context Toolkit

Central to our plans for live use of our system is integrating it into another system called the Context Toolkit. The Context Toolkit aims to ease the development of context aware applications by providing a library of "context widgets" that free the application writer from the details of context sensing (i.e., interfacing with sensors) [13]. In the same way that GUI widgets insulate applications from certain interface presentation concerns, context widgets insulate applications from context acquisition concerns. The system consists of these context widgets and a distributed infrastructure that hosts and coordinates the widgets.

In order to integrate the Smart Floor with the Context Toolkit, we will need to add a software layer that outputs the calculated identity of the user (or perhaps the top three choices, along with a certainty score). We would then have a ready-to-use identity widget, similar to a widget that used another identification technology such as face recognition or RFID tags. Application writers could then easily use this widget as their software interface to the Smart Floor system. They would not need to concern themselves with the details of interfacing to the floor system or with changes to the system as it evolves.

In our current workspace, we are using an application we term the "In/Out Board" which indicates whether specific users are currently in the workspace area [14]. The In/Out Board relies on users to physically check in and out with iButtons (made by Dallas Semiconductor [3]) that they carry with them. The iButton identification system is wrapped into the Context Toolkit, making it much easier to write applications such as the In/Out Board that use the iButton information. With the Smart Floor, we have a system that can replace the iButtons and can provide identity information much more transparently. For example, a Smart Floor tile can be located near the entrance to our workspace and can identify users as they enter and leave. We are currently pursuing this application.

Plate Boundaries

To date, we have implemented the Smart Floor with only a single floor tile. As mentioned above, subjects were instructed to place their foot near the center of the tile. In real use, we cannot instruct users to do this, nor can we be guaranteed of getting a clean footstep every time someone walks over a tile. In the Aware Home, we will have two tiles in a row in some locations (such as the hallway entrance) in order to have a better chance at gathering at least one clean footstep. However, we will need to address the case when the footstep crosses the plate boundaries, or when the user other foot lies partially on one of the plates, artificially loading the plate.

Which Features are Most Important?

In our approach to nearest neighbor classification, we used all ten of the features in calculating the distance between the stored training points and the new, unknown identity point. Further, each feature was given equal weight in the distance measurement. What if it turned our that only four of the ten features were important in this measurement? By equally weighting all ten features, we may actually be increasing the distance between two points in feature space that have very similar values for the four relevant features, and that ought to be classified as the same user.

The solution to this problem is obviously to weight the features differently. We can accomplish this choosing a subset of the training data, and then choosing the feature weights so that the classification error of the remaining training data is minimized [5]. Repeating this process can increase the accuracy of the weights. This algorithm is equivalent to stretching or shrinking each axis of the feature space relative to the others.

CONCLUSION AND SUMMARY

We have designed the Smart Floor system to provide a way to transparently and non-invasively identify users. Our achievable accuracy of over 90% shows that our system is accurate enough for common use. We have also shown that the effect of footwear is negligible on recognition.

We are continuing to investigate and develop appropriate uses and real applications for our system; we are most excited about using the Smart Floor in the Aware Home. This system only scratches the surface of possible applications for this technology, and there are a number of other everyday locations, such as counter tops and chairs, that may benefit from such smart surface technology.

ACKNOWLEDGMENTS

Gary Boone and Sven Koenig gave essential advice about the nearest neighbor algorithms. Irfan Essa, Chris Atkeson, and Thad Starner were very useful in helping us understand HMMs in general and HTK in particular. Thanks to Bob Gregor for the use of his Kistler plate. Jason Turnage has been indispensable in designing and implementing the piezoelectric cable user tracking grid. Gregor Altvater assisted in tracking down references to past work, and in the initial designs of the user tracking grid. Finally, we would like to thank the Graphics, Visualization, and Usability Center at Georgia Tech for supporting this project with a GVU Seed Grant (AY1998-99).

REFERENCES

1. Addlesee, M., Jones, A., Livesey, F., and Samaria, F. The ORL Active Floor. *IEEE Personal Communications*, October 1997, 35-41.
2. Cavanagh, P. The Shoe-Ground Interface in Running. *American Academy of Orthopedic Surgeons Symposium* (Coronado, CA, September, 1980).
3. Dallas Semiconductor. iButton Home Page. Available at http://www.ibutton.com.
4. Kidd, C., Orr, R., Abowd, G., Atkeson, C., Essa, I., MacIntyre, B., Mynatt, E., Starner, T., and Newstetter, W. The Aware Home: A Living Laboratory for Ubiquitous Computing Research. To appear in *Proceedings of CoBuild'99* (Pittsburgh, PA, Oct. 1999).
5. Mitchell, T. *Machine Learning*. WCB/McGraw-Hill, 1997.
6. Measurement Specialties, Inc. Piezo Coax Cable. Available at http://www.msiusa.com/piezo_coax_cable.htm.
7. Orr, R. Smart Floor Home Page. Available at http://www.cc.gatech.edu/fce/smartfloor/index.html.
8. Orr, R., Raymond, R., Berman, J., and Seay, F. A System for Finding Frequently Lost Objects in the Home. *Graphics, Visualization, and Usability Center Technical Report 99-24*, Georgia Tech, 1999. Available at http://www.cc.gatech.edu/gvu/reports/1999/abstracts/99-24.html.
9. Paradiso, J., Abler, C., Hsiao, K., and Reynolds, M. The Magic Carpet: Physical Sensing for Immersive Environments. In *Extended Abstracts of CHI'97* (Atlanta, GA, March 1997), 277-278.
10. Paradiso, J., and Hu, E. Expressive Footwear for Computer-Augmented Dance Performance. In *Proceedings of the First International Symposium on Wearable Computers* (Cambridge, MA.), IEEE Computer Society Press, Oct. 13-14, 1997, 165-166.
11. Pedotti, A. Simple Equipment Used in Clinical Practice for Evaluation of Locomotion. *IEEE Transactions on Biomedical Engineering 24*, 5 (September 1977), 456-461.
12. Pinkston, R., Kerkhoff, J., and McQuiclken, M. A Touch Sensitive Dance Floor/MIDI Controller. In *Proceedings of the 1995 International Computer Music Conference* (Banff, Alta., Sept. 1995), 224-225.
13. Salber, D., Dey, A., and Abowd, G. The Context Toolkit: Aiding the Development of Context-Enabled Applications. In *Proceedings of CHI'99* (Pittsburgh, PA, May 1999), 434-441.
14. Salber, D., Dey, A., Orr, R., and Abowd, G. Designing for Ubiquitous Computing: A Case Study on Context Sensing. *Graphics, Visualization, and Usability Center Technical Report 99-29*, Georgia Tech, 1999. Available at http://www.cc.gatech.edu/gvu/reports/1999/abstracts/99-29.html
15. Santambrogio, G. Procedure for Quantitative Comparison of Ground Reaction Data. *IEEE*

Transactions on Biomedical Engineering 36, 2 (February 1989), 247-255.

16. Want, R., Hooper, A., Falcao, V., and Gibbons, J. The Active Badge Location System. *ACM Transactions on Information Systems 10*, 1, 91-102.
17. Weiser, M. The Computer for the 21st Century. *Scientific American 265* (September 1991), 66-75.
18. Werb, J., and Lanzl, C. Designing a Positioning System for Finding Things and People Indoors. *IEEE Spectrum 35*, 9 (September 1998), 71-78.
19. Wisneski, C., Orbanes, J., and Ishii, H. PingPongPlus: Augmentation and Transformation of Athletic Interpersonal Interaction. In *Extended Abstracts of CHI'98* (Los Angeles, CA, April 1998), 327-328.